

Optimally Fast Incremental Manhattan Plane Embedding and Planar Tight Span Construction

David Eppstein

Computer Science Department
University of California, Irvine
`eppstein@uci.edu`

Abstract. We describe an algorithm for finding the tight span of a finite metric space; the tight span is a construction for embedding arbitrary metrics into L_∞ spaces analogous to the Euclidean convex hull. Our algorithm is incremental, and applies to any space for which the tight span is homeomorphic to a subset of the Euclidean plane. After a new point is added to the metric space our algorithm can update the tight span in time linear in the number of points already added; this is optimal with respect to the size of the input, an $n \times n$ distance matrix. As an application, we improve the running time of an algorithm of Edmonds for embedding finite metrics into the Manhattan-metric plane, from $O(n^2 \log^2 n)$ to $O(n^2)$.

1 Introduction

Unlike the more familiar Euclidean distance, the L_∞ metrics contain all finite metric spaces without distortion: if (X, d) is a finite metric space, each point $x \in X$ can be associated with an $|X|$ -dimensional vector, the distances from x to each point in X , and the L_∞ distance between these vectors equals the original distance between the corresponding points. In contrast, the four-point metric space of distances between the vertices of a claw graph $K_{1,3}$ cannot be embedded into any L_2 space [22].

Edmonds [15] investigated the algorithmic complexity of finding L_∞ embeddings with as low a dimension as possible (for earlier work on this problem and related work on L_1 embedding, see [2–4, 6, 23]). The problem is polynomial for two dimensions [6], but as Edmonds showed it is NP-complete when the dimension is three or more. Even for two-dimensions, L_∞ embedding is significantly more complicated than in the Euclidean plane: in the Euclidean plane, after any two points are placed at the correct distance apart, the position of all remaining points is determined (up to reflection) by their distances from these two points, so one may test in linear time whether a distance matrix represents a planar point set, whereas in the L_∞ (or equivalently L_1) plane, there are infinitely many non-equivalent ways any two points may be placed at a fixed distance from each other, and the placement of two points does not in general determine the location of the remaining points. Despite these complications, Edmonds described an algorithm with running time $O(n^2 \log^2 n)$ for finding an embedding of a given n -point metric space into the Manhattan plane, improving previous bounds [6]; this is nearly optimal, as the input takes the form of a distance matrix with size n^2 .

Low-dimensional L_∞ embeddings have also been studied from the mathematical point of view, using a construction known variously as the *tight span*, *injective envelope*, or *hyperconvex hull* [9–14, 18, 19, 27]. An *injective* or *hyperconvex* metric space [1] shares with L_∞ the properties that it is path-geodesic (any two points may be connected by a line segment) and that its balls have the *Helly property*: any pairwise intersecting family of balls has a point of common intersection. The tight span of a metric space (X, d) is the smallest injective space into which (X, d) can be isometrically embedded: that is, if f is a distance-preserving map from (X, d) into any injective space, f can be extended to a distance-preserving map from the tight span to the same space. Thus, in some sense the tight span is the lowest-dimensional L_∞ -like space into which (X, d) can be embedded, and for points that can be embedded by Edmonds' algorithm into a Manhattan plane, the tight span is essentially the same as the orthogonal convex hull (Figure 1, Lemma 2), a concept well-studied in computational geometry [20, 24–26]. Tight spans arise in phylogenetic analysis as spaces that represent the evolutionary distances between a set of species and are in some sense as tree-like as possible [12, 14]; they have also been used to develop competitive online algorithms for the k -server problem [9]. The tight span of an n -point metric space may be constructed algorithmically, from the subset of bounded faces of a certain n -dimensional unbounded polytope [11, 27], but due to the high dimension of the polytope this method is not computationally efficient and is limited only to spaces with very small numbers of points; in general the tight span may have exponential complexity [18].

We define an injective space to be *planar* if it is homeomorphic to a subset of the plane, and *two-dimensional* if it does not contain any subset homeomorphic to a three-dimensional ball [10, 13]. E.g., three rectangles meeting on a common edge form a space that is two-dimensional but not planar. In prior work [16], we defined and proved injectivity for a class of planar metric spaces that includes but is more general than the Manhattan plane, called *Manhattan orbifolds*. These spaces are topological manifolds in which the metric is locally modeled on the Manhattan plane with the exception of certain *cone points* at which five or more Manhattan-plane quadrants meet. In this paper, we apply Manhattan orbifolds algorithmically to the problems of tight span construction and L_∞ embedding. Specifically, we show the following results:

- Based on our previous work on Manhattan orbifolds, we describe a concrete representation of a planar injective metric space as a *rectangular complex*, a data structure consisting of points, line segments, and L_1 -geometry rectangles.

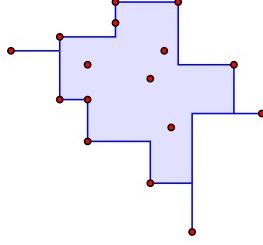


Fig. 1. The orthogonal hull of a set of point sites. For points in the L_1 plane (or, when the points and hull are rotated by 45° , for points in the L_∞ plane), when the hull is connected, it is isometric to the tight span of its sites. [Public-domain figure drawn by the author in 2006 for Wikipedia.]

- We show that, if a metric space with $n - 1$ points has a rectangular complex as its tight span, and if the addition of one more point to the metric space causes its tight span to remain planar, then the tight span of the augmented space can be constructed in $O(n)$ time and that this augmentation causes at most a constant increase in a potential function that counts features of the tight span.
- It follows from the previous result that, whenever a metric space has a planar tight span, the span can be represented in $O(n)$ space and constructed in $O(n^2)$ time, an optimal time bound given the $\Omega(n^2)$ complexity of the input distance matrix. Our algorithm shows that Manhattan orbifolds connected by tree-like features are sufficient to describe all planar tight spans of finite metric spaces.
- We describe a linear time test for whether a rectangular complex can be isometrically embedded into the Manhattan plane. As a consequence, we find an $O(n^2)$ time and $O(n)$ space algorithm for determining whether a finite metric may be embedded into the Manhattan plane, improving on the previous $O(n^2 \log^2 n)$ time algorithm of Edmonds [15].

2 Preliminaries

2.1 Injective metric spaces

A *metric space* is a pair (X, d) , where X is a set of *points* and d is a *distance function*, a function from pairs of points to real numbers satisfying

Positivity. For all x and y , $d(x, y) \geq 0$, with equality only when $x = y$.

Symmetry. For all x and y , $d(x, y) = d(y, x)$.

Triangle inequality. For all x , y , and z , $d(x, y) + d(y, z) \geq d(x, z)$.

An *isometry* between a metric space (X, d) and another metric space (X', d') is a function f from X to X' such that, for all x and y in X , $d(x, y) = d'(f(x), f(y))$. A *geodesic* in a metric space is the isometric image of a line segment; a metric space is *path-geodesic* if every two points are the endpoints of a geodesic.

A family of sets is said to have the *Helly property* if every subfamily that intersects pairwise has a point of common intersection. A *hyperconvex space* is a path-geodesic metric space in which the family of closed balls $B_r(x) = \{y \mid d(x, y) \leq r\}$ has the Helly property. Euclidean spaces of dimension greater than one are not hyperconvex – it is possible to find three planar disks that intersect pairwise but such that all three do not intersect – but L_∞ spaces are hyperconvex. Another familiar example of a hyperconvex space is a tree, represented topologically as a complex of points (vertices of the tree) and line segments (edges of the tree).

Hyperconvexity is equivalent to *injectivity* of a metric space [1]. A *short map* between metric spaces is a function that does not increase the distance between any two points, and a *retraction* is a short map from a metric space to a subspace that restricts to the identity on that subspace. (X, d) is injective if, whenever it forms an isometric subspace of a larger metric space (Y, d) , there exists a retraction from Y to the image of X . In other words, X is an *absolute retract*.

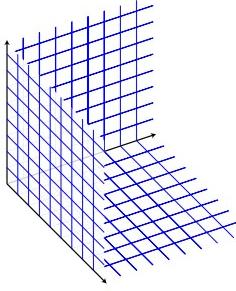


Fig. 2. The order-3 rectilinear cone, formed by gluing together three quadrants of a Manhattan plane. This space is not injective, but analogous spaces formed from larger numbers of quadrants are. From [16].

2.2 The tight span

As Isbell [19] first showed, any metric space (X, d) (finite or infinite) may be embedded isometrically into a unique minimal injective metric space via a construction now known variously as the *tight span*, *injective envelope*, or *hyperconvex hull*. The tight span T_X of a space (X, d) has points that may be represented as the family F of functions from X to \mathbb{R} that satisfy the following properties:

- For each x and y in X , and for each $f \in F$, $f(x) + f(y) \geq d(x, y)$, and
- For each x in X , and for each $f \in F$, $\inf_{y \in X} f(x) + f(y) - d(x, y) = 0$.

We may think of $f(x)$ as the distance from f to x ; with this interpretation, the first property is just the triangle inequality; the first property for $x = y$ forces all function values to be non-negative. The second property forces functions in F to be minimal, in the sense that no function value $f(x)$ can be decreased without violating the triangle inequality. The distance between any two such functions is measured using a form of L_∞ distance:

$$d(f, g) = \sup_{x \in X} |f(x) - g(x)|.$$

The supremum in this formula is bounded, as (choosing $x \in X$ arbitrarily) $d(f, g) \leq f(x) + g(x)$. We may embed (X, d) into F , by mapping any point x to the function $f_x(y) = d(x, y)$. As shown by Isbell [19], this embedding is isometric, and the metric on F produced by this construction is hyperconvex. More, if X has an isometry into any injective metric space (X', d') , then that isometry can be extended to an isometry of T_X into (X', d') , so in some sense T_X is the smallest possible injective metric space into which (X, d) can be isometrically embedded.

We may use this minimality property to identify the tight span of a particular metric space (X, d) : if we exhibit an injective metric space (X', d') into which (X, d) may be isometrically embedded, with the additional properties that each point in X' is uniquely identified by its distances to the points in X and that these distances satisfy the minimality property defining a tight span, then (X', d') is (up to isometry) the tight span of (X, d) . For instance, this argument can be used to show that, when a set of points in the L_1 plane has a connected orthogonal convex hull, this hull is isometric to the tight span (Figure 1, Lemma 2).

2.3 Manhattan orbifolds

The L_∞ plane (or, isometrically, the L_1 plane) forms an example of a two-dimensional injective space, but not the only possible example. In previous work [16] we defined a class of two-dimensional metric spaces which we called *Manhattan orbifolds*, and proved that these spaces are always injective. For completeness, we describe these spaces again here.

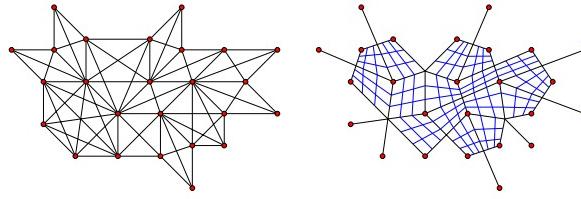


Fig. 3. A non-planar graph and the tight span of its distance metric, a Manhattan orbifold with whiskers. From [16].

Topologically, a Manhattan orbifold is a 2-manifold with boundary: each point has a neighborhood that is homeomorphic to the neighborhood of some point in a plane or half-plane. Every simple closed curve must be the boundary of a unique disk; for instance, this is true for the Euclidean plane but not for the surface of a torus (there exist closed curves that are not boundaries of disks) nor the surface of a sphere (on a sphere, each simple closed curve bounds two disks).

Metrically, a Manhattan orbifold is required to be Cauchy complete (each sequence of points with the property that, for any $\epsilon > 0$, all but finitely many points are at distance ϵ to each other has a limit point) and path-geodesic. Additionally, every point must have a neighborhood of one of five types:

- Ordinary points have a neighborhood isometric to a neighborhood of the origin in the Manhattan plane.
- Boundary-geodesic points have a neighborhood isometric to the neighborhood of a point on the boundary of a region of the Manhattan plane bounded by a smooth curve. Moreover, within this neighborhood the sign of the slope of the curve ($-1, 0, 1$, or ∞) should be constant.
- Inflection points are not ordinary or boundary-geodesic, but have a neighborhood within which all other points are ordinary or boundary-geodesic and that is isometric to the neighborhood of a point in a region bounded by a piecewise smooth curve in the Manhattan plane. At a singular boundary point, the boundary of the manifold may not be smooth, or its image in the Manhattan plane may have a slope that changes sign at that point.
- Cone points of order k have a neighborhood isometric to the neighborhood of the origin in a metric space formed from gluing together k quadrants of the Manhattan plane (Figure 2). In a Manhattan orbifold, all cone points must have order greater than or equal to five.
- Cone inflection points have a neighborhood isometric to the neighborhood of the origin in a region of the order- k rectilinear cone bounded by a smooth curve. All other points in the neighborhood are required to be ordinary or boundary-geodesic.

As we showed in [16], these requirements are enough to ensure that Manhattan orbifolds are injective. Additionally, we used them to identify the tight spans of certain graphs, such as squaregraphs [5, 8] or the graphs formed by replacing the faces of certain planar graphs by cliques (Figure 3).

3 One-point extensions of tight spans

Our algorithm for constructing planar tight spans is *incremental*: it proceeds by adding one point to the input at a time, maintaining as it does the tight span of the points added so far. Therefore, we need some technical results concerning how the tight span may change when a single point is added. Specifically, we prove:

- If S is any subset of a metric space and s any point, the tight span of S has a unique isometric embedding into the tight span of $S \cup \{s\}$ (Lemma 1).
- If P is a geodesic path and s any point, the tight span of $P \cup \{s\}$ can be represented as an orthogonally convex set in the L_1 plane, and additionally if the distances from P to s (as parametrized by path length

along P) are piecewise linear with slopes ± 1 , the tight span of $P \cup \{s\}$ may be partitioned into rectangles only two of which lie along the boundary of the new tight span (Lemma 3).

- If S is any subset of a metric space and s any point, the tight span of $S \cup \{s\}$ attaches to the tight span of S in a path-geodesic subset (Lemma 4).

We defer the details to Appendix I.

4 Rectangular complexes

We define a *rectangular complex* to be a finite set of objects of three types: vertices, edges, and faces. Each face of a rectangular complex represents a metric space with the geometry of an axis-aligned rectangle in the L_1 plane; associated with the face, we store the lengths of the sides of this rectangle, and, for each of the four rectangle sides, a partition of the side into a sequence of edges. Each edge of a rectangular complex represents a metric space with the geometry of a line segment; associated with the edge, we store its length, two vertices at its endpoints, and from zero to two faces that it forms part of the boundary of. We say that an edge is a *bridge* if it is not the boundary of any face. Each vertex of a rectangular complex represents a single point in a metric space; associated with the vertex, we store an unordered list of the edges for which it is the endpoint.

An edge is incident to the faces for which it is part of the boundary, and the vertices at its endpoints; a vertex is incident to a face if there is an edge incident to both the face and the vertex. In order that the rectangular complex have the topology of a subset of the plane, we require that at every vertex the incident faces have one of two types:

- The incident edges and faces may form a single cycle, alternating between edges and faces, such that each two consecutive edges and faces in the cycle are incident. That is, topologically, the vertex has a neighborhood homeomorphic to an open disk in the plane. We define the *angle* of an incident rectangle to be $\pi/2$ if the vertex is one of the four corners of the rectangle, and π if it instead lies on one of the rectangle sides; we require that the total angle of the faces incident with the vertex be at least 2π . That is, in the terminology of Section 2.3, the vertex must be either an ordinary point (with total angle 2π) or a cone point (with total angle greater than 2π).
- Alternatively, the incident edges and faces may form one or more alternating subsequences within which any two consecutive edges and faces are incident, starting and ending with an edge incident to fewer than two faces. If there is exactly one subsequence, which includes at least one face, then the vertex automatically forms a boundary-geodesic point, inflection point, or cone inflection point. However, it is also possible for a vertex to be incident to a single bridge (in which case we call it a *leaf*) or to be incident to more than one subsequence (in which case we call it an *articulation point*).

Additionally, we require that any simple closed curve in the topological space obtained from the complex be the boundary of a unique disk, and that the space be connected.

As an example, the tight span in Figure 3 is not itself a Manhattan orbifold, but it can be represented as a rectangular complex in which the red points and the crossings between black segments are vertices, the portions of the black segments between any two vertices are edges, and the blue crosshatched regions are faces (each representing an L_1 -geometry square with side length $1/2$). Each bridge in the example has a leaf as one of its endpoints and an articulation point as the other. In this example, each of the four sides of a face is formed by a single edge of the complex, but in general a face may have multiple edges and vertices along its sides. The same tight span could also be represented by other rectangular complexes with more or fewer rectangles, by splitting the rectangles into smaller ones or by merging pairs of rectangles into larger ones. In the following section we will discuss additions to the rectangular complex data structure to represent the generating sites of a tight span; we will not require that these sites be vertices of the complex.

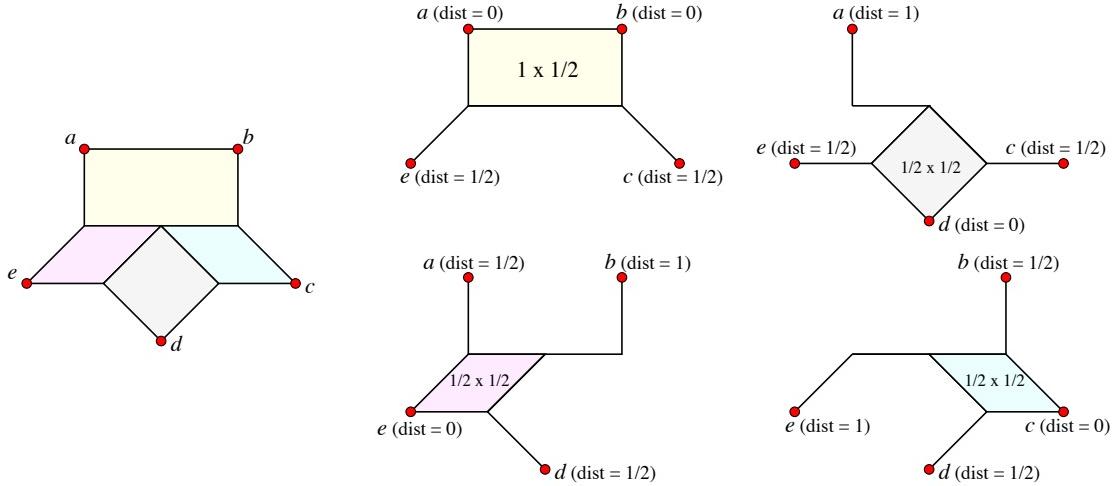


Fig. 4. One possible representation of the tight span of a 5-cycle $abcde$ (with distance 1 between adjacent sites in the cycle and distance 2 between nonadjacent sites) as a rectangular complex, with an exploded view showing the information stored in each face of the complex: its dimensions as an L_1 rectangle, and the distances to four sites that are nearer to each corner of the rectangle than to any other point of the rectangle.

We define a *block* of a rectangular complex to be a maximal set of edges and faces that can be reached from each other through face-edge incidences, together with all of the vertices incident to these edges and faces. It is straightforward to verify that each block must either be a bridge or a Manhattan orbifold, and that the incidences between blocks and articulation points of the complex form a tree. Because bridges (line segments) and Manhattan orbifolds are both injective spaces, and because connecting two injective spaces together by identifying a single point of one with a single point of another preserves injectivity, it follows that every rectangular complex represents an injective metric space.

5 Distances in rectangular complexes

The previous section described rectangular complexes as data structures for representing abstract metric spaces, but we are interested in them as a way to represent more specifically the tight spans of finite metric spaces. In this section we augment the rectangular complex to represent its relationship with the given finite metric space, and we describe how to compute distances in this augmented space. For clarity, we adopt terminology used in computational geometry in the context of Voronoi diagrams: a *site* is one of the points of the finite metric space (X, d) that we are attempting to find the tight span of (or, by extension, its image in the tight span) while a *point* may refer to any point of the rectangular complex under discussion.

Suppose that the rectangular complex R represents the tight span of a finite set (X, d) of sites. We may represent the isometry from (X, d) to R explicitly, by specifying for each site in X which feature of R it maps to and (if that feature is an edge or face) what its position within that feature is. If a site belongs to multiple features, we choose a single one of them arbitrarily.

We also store, as part of our data structure, two sites from X for each bridge of the rectangular complex and four sites from X for each face of the rectangular complex. For each endpoint v of each bridge, we store in the object representing the bridge a pointer to a site $\text{site}(v)$ from X that is closer to v than to any other point of the bridge, and we also store the distance $d(v, \text{site}(v))$. Similarly, for each corner v of each face, we store in the object representing the face a pointer to a site $\text{site}(v)$ from X that is closer to v than to any other point of the face, and we also store the distance $d(v, \text{site}(v))$ (Figure 4). Each bridge or face is therefore a subset of the tight span of the two or four sites associated with that feature.

As we show in Appendix II, we may use this information to calculate the distance between any point of the rectangular complex and any new site, in constant time per calculation.

6 The incremental algorithm

We are now ready to describe our algorithm for adding a site s to a planar tight span described as a rectangular complex C . First, use the distance formula to calculate the distances between s and each feature in C . We verify for each input site t that the distance from s to t equals the distance between s and the point in C representing t ; this can only fail to happen if the input distance matrix does not satisfy the properties of a metric, and in this case we terminate the algorithm with a failure condition.

At any point of C at which the distance to s is a local minimum, the tight span will need to include a new feature that connects s to that point; let S denote the set of local minima found in this way. If S consists of a single point at distance zero from s , then the tight span does not change with the addition of s , and s can be placed at that single point; in this case the algorithm terminates successfully. If any local minimum lies within the interior of a block, the resulting tight span is necessarily nonplanar and the algorithm terminates with a failure condition.

Otherwise, by Lemma 4, the points at which the new tight span features attach to C must contain geodesics between every point in S , and by planarity of the new tight span, the set A of attachment points must form a geodesic path along the boundary of C . Further, if some path follows two adjacent segments around a corner c of some rectangle in C , but c does not belong to S , then this path cannot be the set of attachment points of new features of a tight span, because such an attachment would lead to a cone point of order two or three which is not locally injective. Thus, if the added site leads to a tight span, there must exist a good path in C containing S ; by Lemma 7 there is at most one path P of this type and we can find it efficiently; if P does not exist then the algorithm terminates with a failure condition.

By the distance formulas for distances from s , the distance to S from points on P forms a function that (parametrized by path length) is piecewise linear with slope ± 1 in each piece. Therefore, by Lemma 4, the tight span T of $P \cup \{s\}$ may be represented as a rectangular complex the boundary of which is formed by P together with at most two rectangles and one line segment, with a total number of rectangles equal to $|S| - 1$. We attach T to C to form a new complex $T \cup C$, splitting existing edges of C at points of S when such a point lies in the interior of the edge in order for the result to be a proper complex.

Because $T \cup C$ is formed by attaching a topological disk T to a set C that has the topology of a tree of disks, along a single boundary path, it is homeomorphic to a subset of the plane and satisfies the requirement of a rectangular complex that all simple cycles bound a unique disk. Because P is a good path, and because T contains rectangles with a total angle of $3\pi/2$ at each point of S , there is only one ways that $T \cup C$ can fail to describe a Manhattan orbifold: a point x that is the boundary of a rectangle or rectangles in C that have a total angle of π at that point could form a local maximum of distance from s , so that the angle formed by T at x is $\pi/2$ and the union $T \cup C$ contains a cone point of order three, which is not injective. In this case, if the tight span of $C \cup \{s\}$ were to attach to C via a boundary path through x , a neighborhood of this order-three cone point would be filled in to its injective hull, a three-dimensional octant of the three-dimensional L_∞ space, violating planarity. Thus, we may test whether $T \cup C$ describes a Manhattan orbifold by examining the total angle at each of its vertices, and if any of these local angles is too small we terminate the algorithm with a failure condition.

If all of these steps succeed, then $T \cup C$ is an injective space that contains all the points. Every point in $T \cup C$ either belongs to the tight span of the previously added points, or the tight span of $P \cup \{s\}$, so $T \cup C$ must equal the tight span of the point set with s included.

It remains to determine the two representative sites for any new bridge and the four representative sites for any new rectangular complex feature created by this algorithm, in order to be able to apply the distance formula in subsequent iterations. In the case of a bridge, one representative site is s and the other can be any

previously added site. In the case of a rectangle, three representative sites can be taken to be s and the two endpoints of P ; the fourth can be taken as any existing site that determines the distance from s according to the distance formula at the fourth corner of the rectangle, a point of P at which the distance from s is a local maximum. Once these sites are determined for each feature, the algorithm terminates successfully.

Theorem 1. *Suppose that a finite metric space (X, d) has a planar tight span. Then this tight span can be represented as a rectangular complex with $O(n)$ features, and can be constructed by applying the algorithm described above to the points, one at a time in any order. The time to add each point is $O(n)$ and the total time to compute the overall tight span is $O(n^2)$.*

Proof. The correctness of the algorithm follows from the arguments sketched above: the algorithm only terminates with a failure condition when some subset of X has a tight span that is guaranteed to be nonplanar, and if it terminates successfully then it has correctly found the tight span of the subset of points considered so far.

To prove the asymptotic bounds on the number of features of the rectangular complex and the running time of the algorithm, consider the potential function $\Phi = R + 3B + E$ where R is the number of rectangles in the current rectangular complex, B is the number of bridges, and E is the number of edges that are on the boundary of a block. Whenever a point is added, with a set S of local minima of the distance function, the addition creates $|S| - 1$ new rectangles and at most one new bridge, increasing Φ by $|S| + 2$. In addition, if the edges at either end of P were bridges, they may be split into a bridge and a boundary edge, increasing Φ by two more units total. However, at least $|S| - 2$ edges of the previous rectangular complex (an edge containing or incident to each local minimum in S ; note that no two local minima can share an edge because the local maximum between them would lead to a nonplanarity) lie entirely along the sides of newly added rectangles. If one of these edges incident to an interior local minimum was a bridge, it becomes one or two boundary edges, and if it was a boundary edge, it becomes an interior edge of its block, so these changes reduce Φ by at least $|S| - 2$. Thus, in all cases, the total change to Φ per point of X is at most an increase by six, and after the algorithm completes $\Phi \leq 6|X|$. This bounds the number of rectangles and bridges of the rectangular complex by $6n$. By a standard computation involving the Euler formula, the number of edges can be at most $18n$ (in any planar subdivision, the number of adjacencies between faces is at most three times the number of faces) and the number of vertices can be at most $12n$ for similar reasons.

Since each computation takes time linear in the size of the subdivision so far together with the number of newly added features, and we have seen that the size of the subdivision remains linear throughout the algorithm, the total time per point is linear and the total time for the whole computation is quadratic. \square

7 Manhattan plane embedding

By injectivity, if a metric space can be embedded isometrically into the Manhattan plane, then so can its tight span. Therefore, we can find Manhattan plane embeddings by using the algorithm of the previous section to find a planar tight span of the given metric space, and using the structure of the tight span to determine its possible Manhattan plane embeddings. We remark that, for metrics that can be embedded into the Manhattan plane, the constant factors in the analysis of Theorem 1 can be significantly tightened: the interior local minima of distance on the path connecting each new part of the tight span to the previously constructed part must each be sites, so if the algorithm creates k new rectangles then it removes $k - 1$ sites from the boundary of the tight span. The first rectangle cannot be created until there are four sites, and (if any rectangles are created) at least four sites remain on the boundary at the end of the algorithm. Thus, in this case, the total number of rectangles formed for a metric space with n sites is at most $2n - 7$. This bound is tight: there exist metric spaces formed by points in the Manhattan plane, and insertion orderings for those points, that cause our algorithm to form $2n - 7$ rectangles (Figure 5).

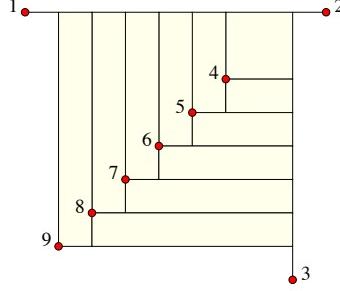


Fig. 5. An insertion ordering for a set of sites in the L_1 plane that leads to a rectangular complex with $2n - 7$ rectangles.

We define a *hinge* of a rectangular complex to be an articulation point of the complex that connects more than two blocks, or that connects exactly two blocks and has the additional property that neither of the two subcomplexes formed by cutting the complex at the articulation point has the structure of a path of bridges.

As we now discuss, there are several straightforward necessary conditions describing a rectangular complex that comes from the distances of a set of point sites in the Manhattan plane. The justification for these conditions will assume a fixed embedding of the sites into the plane; this also determines the locations of the points of the complex that belong to a block of the complex that is not a bridge. The points interior to a bridge do not necessarily have a unique representation as a point in the Manhattan plane; however, whenever our algorithm creates or modifies a bridge, each endpoint of the bridge is either a site or a point of a block that contains a rectangle, so the positions of the two bridge endpoints can also be uniquely determined as points in the plane from the positions of the sites. With this determination of locations for points of the complex, the complex may be seen to satisfy the following conditions:

- Any two blocks of the complex that meet at an articulation point must come from points that are separated geometrically by an empty open quadrant of the plane. Thus, at any articulation point, there must be at least two diagonally opposite empty open quadrants.
- Each articulation point of the complex must be incident to at most four blocks. If it is incident to exactly four blocks, the complex must have the structure of a tree with one degree-four vertex connecting four paths of bridges and articulation points. For, a point set that lies entirely on one axis-aligned line will lead to a rectangular complex that is a path, so any complex that is not a path must extend over a wider angle, and having four such complexes would violate the empty quadrant property described above.
- For the same reason, if an articulation point is incident to three blocks, at least two of them must form paths of bridges and articulation points.
- Within a block that is not a bridge, each boundary vertex is surrounded by rectangles forming a total angle of at most $3\pi/2$ with that vertex, and each non-boundary vertex is surrounded by rectangles forming a total angle of exactly 2π . For, otherwise the geometry in a neighborhood of that vertex would be that of a higher-order cone point, not possible in the Manhattan plane.
- Within a block that is not a bridge, define a *boundary segment* to be a maximal path of consecutive boundary edges and vertices with the property that the angle at each internal vertex of the path is exactly π . Then there must be exactly four *extremal segments* for which the angles at the vertices at each end of the segment are $\pi/2$; all remaining boundary segments must have one end vertex with angle $\pi/2$ and the other with angle $3\pi/2$. For otherwise, the block could not form an orthogonally convex simple polygon in the plane.
- Within a block that is not a bridge, each articulation point must lie on an extremal segment. For, at the other boundary points of the block, the block extends into three of the four quadrants, making it impossible to find two opposite empty quadrants for the articulation point.

- If B is a block that is not a bridge, every hinge of B must coincide with the vertex where two extremal segments meet. For again, otherwise the condition of having two opposite empty quadrants could not be met.
- Within any block B that is not a bridge, there must exist an assignment of articulation points to extremal segments of B , such that each hinge is assigned to both of its adjacent extremal segments, each non-hinge articulation point is assigned to one extremal segment, and each extremal segment has at most one articulation point assigned to it. This assignment may be determined by letting the assigned edge of a path be the extremal segment that is embedded perpendicularly to the path edge incident to it at the articulation point. No two articulation points can be assigned in this way to the same extremal segment because then the condition of having two opposite empty quadrants would be violated at both of them.

As we show in Appendix III, these conditions are necessary and sufficient for a rectangular complex to be embedded isometrically into the Manhattan plane, and may be tested in linear time. As a consequence, we have:

Theorem 2. *Given any rectangular complex, we may test whether it can be embedded isometrically into the Manhattan plane, and if so find an embedding, in linear time.*

Theorem 3. *If we are given as input a finite metric space (X, d) with n points, we may determine whether it represents the distances between n points of the Manhattan plane in time $O(n^2)$.*

Proof. Apply Theorem 1 to compute the tight span C of (X, d) , and apply Theorem 2 to test the embeddability of C and find an embedding for C . \square

8 Discussion

We have characterized the metric spaces that may be formed as a planar tight span of a finite metric space, shown that they may be represented concisely and that distances may be computed quickly from this concise representation, and used these results to develop an efficient algorithm for constructing planar tight spans and for finding embeddings of metric spaces into the Manhattan plane.

Although it is known that a metric space can be embedded into the Manhattan plane if and only if this is true of every six-point subspace [3] we observe that no such finite criterion exists for having a planar tight span. For, let G be a graph formed from the disjoint union of a k -cycle (with $k \geq 4$) and a single isolated vertex, and form a metric space on the vertices of G with distance one between adjacent vertices and two between nonadjacent vertices. Then G has a nonplanar tight span in the form of $k L_1$ squares of side length $1/2$ connected at an order- k cone point, together with a bridge leading from that cone point to the isolated vertex. However, removing any one point from G leads to a metric space with a planar tight span.

The next simplest case to investigate would seem to be the metric spaces having two-dimensional tight spans. Can these tight spans be represented concisely and constructed efficiently? As was already shown by Dress [11], the tight span of five points is always two-dimensional, but if it is not planar then it will contain internal boundaries that are not aligned with the axes of an L_1 -plane representation of its local features, so in order to handle this case it would be necessary to develop a more general representation of two-dimensional tight spans that is not based purely on rectangles; see [21] and [7] for additional mathematical investigations of two-dimensional tight spans.

More ambitiously, it would be of interest to determine for any fixed d the complexity of representing and constructing a tight span of a finite metric space that is guaranteed to be at most d -dimensional. The existence of polynomial time algorithms for this problem, for any fixed d , is not ruled out by Edmonds' NP-completeness proof [15]; rather, that proof shows that it is likely to be hard to determine whether a tight span of dimension at most three can be embedded isometrically into L_∞^3 .

References

1. N. Aronszajn and P. Panitchpakdi. Extensions of uniformly continuous transformations and hyperconvex metric spaces. *Pacific Journal of Mathematics*, 6:405–439, 1956.
2. M. Bâdoiu. Approximation algorithm for embedding metrics into a two-dimensional space. In *Proc. 14th ACM-SIAM Symp. Discrete Algorithms (SODA 2003)*, pages 434–443, 2003.
3. H.-J. Bandelt and V. Chepoi. Embedding metric spaces in the rectilinear plane: a six-point criterion. *Discrete and Computational Geometry*, 15(1):107–117, 1996.
4. H.-J. Bandelt and V. Chepoi. Embedding into the rectilinear grid. *Networks*, 32(2):127–132, 1998.
5. H.-J. Bandelt, V. Chepoi, and D. Eppstein. Combinatorics and geometry of finite and infinite squaregraphs. In preparation, 2009.
6. H.-J. Bandelt, V. Chepoi, and M. Laurent. Embedding into rectilinear spaces. *Discrete and Computational Geometry*, 19(4):595–604, 1998.
7. V. Chepoi. Graphs of some CAT(0) complexes. *Advances in Applied Mathematics*, 24(2):125–179, 2000.
8. V. Chepoi, F. Dragan, and Y. Vaxès. Center and diameter problems in plane triangulations and quadrangulations. In *Proc. 13th ACM-SIAM Symp. Discrete Algorithms (SODA 2002)*, pages 346–355, 2002.
9. M. Chrobak and L. L. Larmore. Generosity helps, or an 11-competitive algorithm for three servers. *Journal of Algorithms*, 16:234–263, 1994.
10. M. Develin. Dimensions of tight spans. *Annals of Combinatorics*, 10(1):53–61, 2006.
11. A. W. M. Dress. Trees, tight extensions of metric spaces, and the cohomological dimension of certain groups. *Advances in Mathematics*, 53:321–402, 1984.
12. A. W. M. Dress, K. T. Huber, and V. Moulton. Metric spaces in pure and applied mathematics. In *Proceedings Quadratic Forms LSU*, Documenta Mathematica, pages 121–139, 2001.
13. A. W. M. Dress, K. T. Huber, and V. Moulton. Totally split-decomposable metrics of combinatorial dimension two. *Annals of Combinatorics*, 5(1):99–112, 2001.
14. A. W. M. Dress, K. T. Huber, and V. Moulton. An explicit computation of the injective hull of certain finite metric spaces in terms of their associated Buneman complex. *Advances in Mathematics*, 168(1):1–28, 2002.
15. J. Edmonds. Embedding into l_∞^2 is easy, embedding into l_∞^3 is NP-complete. *Discrete and Computational Geometry*, 39(4):747–765, 2008.
16. D. Eppstein. Manhattan orbifolds. Electronic preprint arxiv:math/0612109, 2006.
17. A. V. Goldberg and C. Harrelson. Computing the shortest path: A* search meets graph theory. In *Proc. 16th ACM-SIAM Symp. Discrete Algorithms (SODA 2005)*, pages 156–165, 2005.
18. S. Herrmann and M. Joswig. Bounds on the f -vectors of tight spans. Electronic preprint arxiv:math.MG/0605401, 2006.
19. J. R. Isbell. Six theorems about injective metric spaces. *Comment. Math. Helv.*, 39:65–76, 1964.
20. R. G. Karlsson and M. H. Overmars. Scanline algorithms on a grid. *BIT*, 28(2):227–241, 1988.
21. A. V. Karzanov. Metrics with finite sets of primitive extensions. *Annals of Combinatorics*, 2(3):211–241, 1998.
22. N. Linial. Finite metric spaces—combinatorics, geometry and algorithms. In *Proc. International Congress of Mathematicians, Beijing*, volume 3, pages 573–586, 2002.
23. S. M. Malitz and J. I. Malitz. A bounded compactness theorem for L^1 -embeddability of metric spaces in the plane. *Discrete and Computational Geometry*, 8(1):373–385, 1992.
24. D. Y. Montuno and A. Fournier. Finding the x - y convex hull of a set of x - y polygons. Technical Report 148, University of Toronto, 1982.
25. T. M. Nicholl, D. T. Lee, Y. Z. Liao, and C. K. Wong. On the X - Y convex hull of a set of X - Y polygons. *BIT*, 23:456–471, 1983.
26. T. Ottman, E. Soisalon-Soisinen, and D. Wood. On the definition and computation of rectilinear convex hulls. *Information Sciences*, 33:157–171, 1984.
27. B. Sturmfels and J. Yu. Classification of six-point metrics. *Electronic Journal of Combinatorics*, 11:R44, 2004.

Appendix I: One-point extensions of tight spans

As our algorithm depends on adding one point to a tight span, we need some technical results describing how the addition of a point can change the tight span.

Lemma 1. *Let (X, d) be any metric space, and $Y = X \setminus \{y\}$. Then there exists a unique isometric embedding from the tight span T_Y of Y to the tight span T_X of X that maps Y to the image of Y in T_X .*

Proof. Represent each point of T_Y , as above, as a function f from Y to \mathbb{R} satisfying certain constraints. Then there is a unique way to augment this function to one from X to \mathbb{R} satisfying the same constraints: we must set $f(y) = \sup_{x \in X} d(x, y) - f(x)$. \square

This unique embedding property depends on adding only one point at a time. For instance, if four points $abcd$ form a rectangle in the L_1 plane, their tight span is that rectangle, but the tight span of two opposite corners of the rectangle (a line segment between the two) may be embedded isometrically into the rectangle as any of infinitely many different monotone paths.

In order to prove our next extension property, we need to study injectivity in the Manhattan plane. We say that a point h in the L_1 plane is *surrounded* by X if each of the four closed axis-aligned quadrants centered at h contains at least one point of X , and we define the convex hull H of X to consist of all points surrounded by X (Figure 1).

Lemma 2. *For any subset X of the L_1 plane, if the orthogonal convex hull H of X is connected, it is isometric to the tight span of X .*

Proof. By injectivity, the tight span of X is isometric to a subset of the L_1 plane; that is, if we identify the tight span with the set of vectors describing distances from points in the span to sites in X , we need only consider distance vectors determined in this way from points in the L_1 plane. No two points within the convex hull determine the same distance vector. If a point p in H is surrounded by X , its vector of distances to X satisfies the requirements of the tight span that, for each site q in X there exists an r such that $d(p, q) + d(p, r) = d(q, r)$: namely, take r to be a point in the opposite quadrant from q . If a point p is not surrounded, that is, it has an empty quadrant, and if the interior of the opposite quadrant is nonempty, then this requirement cannot be satisfied for those points in the nonempty quadrant, so the distance function for p does not belong to the tight span. And if a point p is not surrounded, but has two diagonally opposite empty quadrants and two nonempty quadrants, then the convex hull is a subset of the interiors of the two nonempty quadrants and is not connected. Thus, in each case it can be shown that either the convex hull coincides with the tight span or the convex hull is not connected. \square

If P is a geodesic path in a metric space, then it is injective and is therefore its own tight span. We next investigate the case when we add one point to P . The result depends on the shape of the function that maps points in P to their distance to the new point; we define a *sawtooth function* to be a piecewise linear univariate function that, in any linear piece of the function, has slope ± 1 , and that has finitely many breakpoints.

Lemma 3. *Let (X, d) be a metric space, P a geodesic path in X , and s a point in X that is not on P . Then the tight span of $P \cup \{s\}$ can be isometrically embedded into an orthogonally convex subset of the L_1 plane bounded between three geodesic paths, one of which is the image of P and the other two of which connect s to the endpoints of P .*

If, in addition, the function that maps points of P (parametrized from their distance from one endpoint of P) to their distance to the new point is a sawtooth function, the tight span can be partitioned into a number of L_1 rectangles equal to one less than the number of local minima of the sawtooth function, together with P and possibly a single line segment, such that the boundary of the tight span is covered by P , by two of the rectangles, and by the line segment.

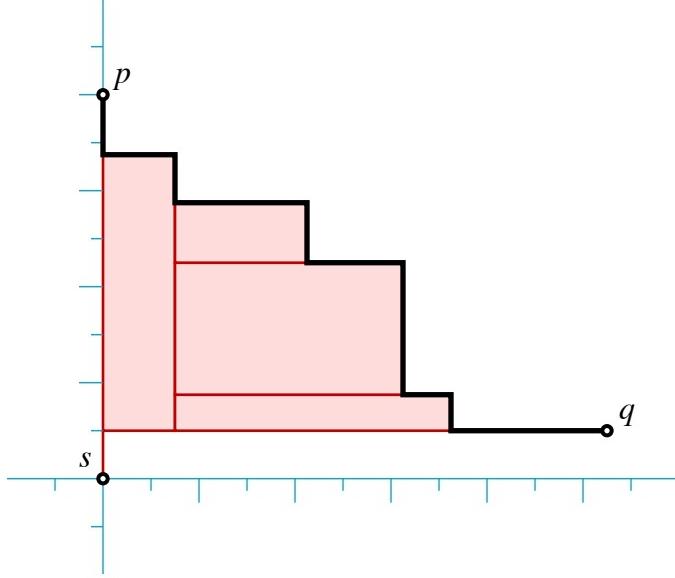


Fig. 6. Illustration for Lemma 3, showing the embedding of path pq in the upper right quadrant of the L_1 plane (heavy black edges), the additional point s at the origin, and the tight span formed as the convex hull of the path and s (medium red). In this example, the distances from the path to s form a sawtooth function and the tight span has accordingly been partitioned into rectangles, such that only two rectangles lie on the boundary of the span.

Proof. Let the endpoints of P be the points p and q , and map each point r of P to the point

$$(x_r, y_r) = \left(\frac{d(p, r) + d(r, s) - d(p, s)}{2}, \frac{d(p, s) - d(p, r) + d(r, s)}{2} \right)$$

of the L_1 plane, as depicted in Figure 6. It follows from the triangle inequality that this point lies within the positive quadrant of the plane, and that the distance between the images of any two points in P is equal to their distance in P , so this is an isometric embedding. The isometry may be extended to $P \cup \{s\}$ by mapping s to the origin. In the case that the function from P to its distances from s is a sawtooth function, the image of P is a polygonal chain in which each edge is axis-parallel. Now let T be the orthogonal convex hull of the image of P and of the origin. T is connected: it has the form of an orthogonal polygon bounded to the left by the vertical line segment ps along the y axis, from above and to the right by P itself, and from below and to the right by a path from s to q that follows the y axis and the horizontal line through q . Therefore, by Lemma 2, it is the tight span. In the case of a sawtooth function, we may cover T with rectangles as shown in the figure. \square

If (X, d) is a metric space, $y \in X$, and $Y = X \setminus \{x\}$, let T_X and T_Y be the tight spans of X and Y respectively. We define a point $t \in T_Y$ to be a *weak attachment point* for y if every neighborhood of t contains a point in $T_X \setminus T_Y$, and a *strong attachment point* if there is no geodesic from t to y that passes through any other point of T_Y . Observe that a strong attachment point is automatically a weak attachment point.

Lemma 4. *Let $X = Y \cup \{y\}$, and let s and t be strong attachment points for y in T_Y . Then there exists a geodesic from s to t in T_Y in which all points are weak attachment points.*

Proof. Let R be a maximal set of strong attachment points that all lie on a single geodesic P from s to t in T_Y ; then P may be partitioned into intervals within which the strong attachment points are dense and

intervals within which there are no strong attachment points. The set of weak attachment points is closed, so P passes through only weak attachment points within each dense region; we need only worry about the intervals within which there are no strong attachment points. Therefore, assume without loss of generality that s and t bound a single such interval; that is, that there are no strong attachment points on any geodesic from s to t .

Now let T_{Py} be the tight span of $P \cup \{y\}$, as described by Lemma 3; T_{Py} contains some points of T_Y (including all points in P) and other points that do not belong to T_Y . Let R be the bounding rectangle of T_{Py} . If the interior of R contains any points of T_Y , let r be such a point that is as close as possible to y ; then r must be a strong attachment point, contradicting the assumption that there are none on any geodesic from s to t . If, on the other hand, all points in $T_{Py} \cap T_Y$ lie on the boundary of R , then in particular the path P follows this boundary and every point in P is a weak attachment point. \square

Appendix II: Distance formulas

Recall that our representation of a rectangular complex stores, with each bridge or face of the complex, two or four sites (respectively) that are closer to each bridge endpoint or cell corner than to any other point within the bridge or face. From this information, we may calculate the distance between any point of the rectangular complex and any new site, using only information that can be looked up in constant time, via a formula resembling a landmark-based technique of Goldberg and Harrelson for bounding distances in arbitrary metric spaces [17]:

Lemma 5. *Let C be a rectangular complex, augmented as above, that represents the tight span of a set S of $n - 1$ sites, let R be a rectangular face of C , let r be a point in R , and let s be an additional site. Let the four corners of R be the points v_0, v_1, v_2 , and v_3 . Then, in the tight span of $S \cup \{s\}$,*

$$d(r, s) = \max_i d(s, \text{site}(v_i)) - d(v_i, \text{site}(v_i)) - d(r, v_i).$$

Proof. By Lemma 1, $d(r, s)$ is well-defined. Let T be the tight span of $\{s, v_0, v_1, v_2, v_3, v_4\}$. By Lemma 1, the tight span of the subset $\{v_0, v_1, v_2, v_3, v_4\}$, and hence the rectangle R , has a unique isometric embedding into T . The same rectangle R , defined as a subspace of T , has a unique embedding into C that respects the distances to $\{v_0, v_1, v_2, v_3, v_4\}$. By injectivity, T is an isometric subset of the tight span of $S \cup \{s\}$, so the distance between r and s in T equals that in the tight span of $S \cup \{s\}$. But in T , the requirements that for each i , $d(r, s) \geq d(s, \text{site}(v_i)) - d(v_i, \text{site}(v_i)) - d(r, v_i)$ and that there exists some i such that this inequality is an equality are just instances of the two properties defining the tight span as a subset of the functions from $\{s, v_0, v_1, v_2, v_3, v_4\}$ to \mathbb{R} . \square

The same proof shows an analogous result for the bridges of a rectangular complex.

Lemma 6. *Let C be a rectangular complex, augmented as above, that represents the tight span of a set S of $n - 1$ sites, let B be a bridge of C , let r be a point in B , and let s be an additional site. Let the two endpoints of B be v_0 and v_1 . Then, in the tight span of $S \cup \{s\}$, $d(r, s) = \max_i d(r, \text{site}(v_i)) - d(v_i, \text{site}(v_i)) - d(r, v_i)$.*

We also need a technical result describing geodesics that follow the boundary of a rectangular complex. Given a subset S in a rectangular complex, define a *good path* for S to be a single geodesic that contains all points of S , contains only points that are on the boundary of C , and does not contain points from two adjacent sides of a rectangle in the complex unless the corner of the rectangle belongs to S .

Lemma 7. *Let S be any subset of boundary points of a rectangular complex C . Then there is at most one good path for S in C . We can test whether this path exists, and if so construct it, in time linear in the number of objects in C .*

Proof. If S spans multiple bridges and blocks of C , the geodesic must follow a path in the bridge-block tree of C , and must include all articulation points of this path. Thus, since the case for a single bridge is obvious, we can reduce the problem to the case in which C consists of a single block, which topologically must have the form of a disk. Let s be any point in S , and consider the distances from s around the boundary of this disk. At s itself, the distance from s is zero, and it increases in both directions as one moves along the boundary away from s . If it decreases again to a local minimum, the geodesic cannot pass through this region; the remaining portion of boundary has the topology of an interval of a line and it is straightforward to find the necessary path. Otherwise, the distances from s increase in both directions to a single local maximum opposite s on the boundary of C , and the only case of ambiguity arises when s and the local maximum are the only points of S on C . In this case, it follows from a lemma in [16] that C has at least four corners of single rectangles on its boundary, at most two of which can be s and the local maximum, so only one of the two paths from s to the local maximum can avoid passing through the adjacent sides to other two corners.

Algorithmically we may pick a single point $s \in S$, use the previous lemmas to compute the distances of all other points on the boundary of C from s , and use the reasoning above to determine a path following the boundary that must be the desired path, if the desired path exists at all. We may test whether the path found by this process is a geodesic by computing its length (the sum of the lengths of its segments) and using the distance formula to check whether its endpoints are that length apart. \square

Appendix III: Characterization of Manhattan-plane embeddability

Section 7 listed a set of conditions that are necessary for any rectangular complex to be embedded into the Manhattan plane. As we now show, these conditions are also sufficient for such an embedding to exist. For convenience, we repeat the conditions here.

- Any two blocks of the complex that meet at an articulation point must come from points that are separated geometrically by an empty open quadrant of the plane. Thus, at any articulation point, there must be at least two diagonally opposite empty open quadrants.
- Each articulation point of the complex must be incident to at most four blocks. If it is incident to exactly four blocks, the complex must have the structure of a tree with one degree-four vertex connecting four paths of bridges and articulation points. For, a point set that lies entirely on one axis-aligned line will lead to a rectangular complex that is a path, so any complex that is not a path must extend over a wider angle, and having four such complexes would violate the empty quadrant property described above.
- For the same reason, if an articulation point is incident to three blocks, at least two of them must form paths of bridges and articulation points.
- Within a block that is not a bridge, each boundary vertex is surrounded by rectangles forming a total angle of at most $3\pi/2$ with that vertex, and each non-boundary vertex is surrounded by rectangles forming a total angle of exactly 2π . For, otherwise the geometry in a neighborhood of that vertex would be that of a higher-order cone point, not possible in the Manhattan plane.
- Within a block that is not a bridge, define a *boundary segment* to be a maximal path of consecutive boundary edges and vertices with the property that the angle at each internal vertex of the path is exactly π . Then there must be exactly four *extremal segments* for which the angles at the vertices at each end of the segment are $\pi/2$; all remaining boundary segments must have one end vertex with angle $\pi/2$ and the other with angle $3\pi/2$. For otherwise, the block could not form an orthogonally convex simple polygon in the plane.
- Within a block that is not a bridge, each articulation point must lie on an extremal segment. For, at the other boundary points of the block, the block extends into three of the four quadrants, making it impossible to find two opposite empty quadrants for the articulation point.

- If B is a block that is not a bridge, every hinge of B must coincide with the vertex where two extremal segments meet. For again, otherwise the condition of having two opposite empty quadrants could not be met.
- Within any block B that is not a bridge, there must exist an assignment of articulation points to extremal segments of B , such that each hinge is assigned to both of its adjacent extremal segments, each non-hinge articulation point is assigned to one extremal segment, and each extremal segment has at most one articulation point assigned to it. This assignment may be determined by letting the assigned edge of a path be the extremal segment that is embedded perpendicularly to the path edge incident to it at the articulation point. No two articulation points can be assigned in this way to the same extremal segment because then the condition of having two opposite empty quadrants would be violated at both of them.

Lemma 8. *It is possible to test each of the conditions above, except for the first one concerning pairs of empty open quadrants, in time linear in the number of features of a given rectangular complex.*

Proof. The remaining conditions are combinatorial in nature and do not depend on a fixed embedding of the complex, so they may be tested without reference to an embedding. A standard connectivity algorithm (in a dual graph formed by rectangles and edges) allows the complex to be partitioned into its blocks in linear time, after which we may calculate the angle formed at each vertex within each block. By searching from leaves inwards it is possible to identify each bridge of the complex that leads to a path of bridges and articulation points, and thereby determine which articulation points are hinges. Most of the remaining conditions are local and concern only this information about total angle and whether an articulation point leads to a path, and may be tested in constant time at each articulation point. The identification of boundary segments and extremal segments is straightforward; if a block has more than four articulation points then it cannot be possible to assign articulation points to extremal segments as described by the final condition, while otherwise there are finitely many possible assignments to test and the assignment can be found in constant time per block. \square

Lemma 9. *If a rectangular complex R satisfies all the conditions of Lemma 8, the tight span of the hinges in the complex forms a path.*

Proof. Assume for a contradiction that the tight span of the hinges does not form a path, then some triple of hinges h_0, h_1 , and h_2 has a tight span T that is not a path. By injectivity, T may be embedded isometrically into R . The tight span of three points is either a path or three paths connected at one central point; let p be an image of this central point in R under the isometric embedding of T into R . If p is not itself an articulation point of R , then the block containing p has at least three incident hinges (one that either coincides with or leads to each hinge h_i) contradicting the assignment of incident hinges to pairs of extremal segments. If on the other hand p is an articulation point of R , then it is incident to three non-path subsets of R , violating the third bullet point above. This contradiction completes the proof. \square

Lemma 10. *Each block that is not a bridge in a rectangular complex C that satisfies the conditions of Lemma 8 has an embedding into the L_1 plane as an orthogonal polygon, with the extremal segments lying along the bounding box of the polygon. The embedding may be found in linear time and is unique up to translations, rotations by angles that are multiples of $\pi/2$, and reflections of the plane.*

Proof. Choose arbitrarily one point o of C to be placed at the origin of the embedding and one orientation for the complex at that point. For each other point p of C we may determine its embedding in the Manhattan plane by tracing a curve in C from o to p and tracing a matching curve in the plane from the origin to the image of p ; the choice of curve is irrelevant, because any two curves from o to p in C can be concatenated together to form a simple closed curve in C , and the shortest two curves that lead to an inconsistent placement would form a curve that does not bound a disk in C , contradicting the requirement that every simple closed

curve in a rectangular complex bounds a disk. In particular, if p and q are any two points, a curve from o to p and then via a geodesic to q is mapped in this way to a curve in the plane that includes an equal-length geodesic from p to q , showing that this embedding is isometric. Every rectangle is thus embedded in an axis-aligned way, so the boundary of the block must form an orthogonal polygon; it must be an orthogonally convex polygon because otherwise it would not be injective, from which the placement of the extremal segments on the bounding box follows. \square

Lemma 11. *If a rectangular complex satisfies all the conditions of Lemma 8 then it can be embedded isometrically into the Manhattan plane; an embedding of this type can be found in time linear in the size of the complex.*

Proof. We embed the blocks of the complex in the order given by the path of hinges from Lemma 9, maintaining an invariant that, at each step, the blocks on one side of a hinge have been embedded within the closed negative quadrant with respect to that hinge.

To begin the embedding, suppose that there exists a hinge, and that the first hinge on the path of hinges is a hinge incident to two paths of bridges and articulation points and to a third block. We then embed the hinge at the origin of the plane and the two paths along the two negative coordinate axes. Alternatively, suppose that the first hinge of the path of hinges connects to a block B that is not a bridge and that has no other incident hinge; then B may be embedded in the plane by Lemma 10 and oriented in such a way that the hinge of B is the upper right vertex of its bounding box; any path connected to B by a non-hinge articulation point may be embedded along a line perpendicular to the extremal segment assigned to that articulation point.

Once we have embedded part of the complex in this way, if a block connecting two hinges is a bridge then the next hinge may be embedded at any point the correct distance from the previous hinge within the positive quadrant relative to the previous hinge. If the block is not a bridge, an embedding of it may be given by Lemma 10, in such a way that the previous hinge is the lower left corner of its bounding box and the next hinge is the upper right corner; there can be no other articulation points of the block. In this way the invariant that the embedded portions of the complex lie to the lower left of the last hinge is preserved.

When the final hinge of the path of hinges is reached, it may connect to two paths, which may be embedded on lines extending upwards and rightwards of the hinge. Alternatively, it may connect to a block with a single hinge; an embedding of this block may be given by Lemma 10, placed in such a way that the hinge lies on the lower left corner of the bounding box of the block, and again any path connected to the block by a non-hinge articulation point may be embedded along a line perpendicular to the extremal segment assigned to that articulation point.

It remains to consider complexes that have no hinges. Such a complex may consist of a single path, in which case it may be embedded along a coordinate axis of the plane. Alternatively, it may have a single non-bridge block, which may be embedded by Lemma 10, after which any path connected to the block by a non-hinge articulation point may be embedded along a line perpendicular to the extremal segment assigned to that articulation point. \square